Q&A with Capt. John D. Zimmerman
Deputy CIO Naval Sea Systems Command
Honor, Courage, Commitment and Cyber

By CHIPS Magazine - April-June 2016 **

In a recent article, "Thinking Slow on Cybersecurity," Capt. Zimmerman, the Deputy Command Information Officer at Naval Sea Systems Command, challenged some of the cybersecurity concepts and best practices used by the Department of Defense and the Navy to secure networks and information systems as an unproductive and costly response to threats and vulnerabilities. In this interview he addresses cyber value, strategy, certification and accreditation, red teams, and cyber ranges — as well as the thinking and values he believes are critical to effectively address our cybersecurity challenges.

CHIPS asked Capt. Zimmerman to share his thoughts on the cybersecurity dilemma, starting with his question, "What better technical, operational, and process solutions can deliver us the most cyber value, in the least amount of time — at the least cost?"

Q: I think your question sums up the cybersecurity issue in a nutshell. A lot of organizations just throw money at the problem. How do organizations get the most cyber value, in the least amount of time — at the least cost?

A: That is the billion-dollar question. First, thanks for having me again. I'm glad you have invited me back and I look forward to discussing a number of important cybersecurity issues with you.

As Navy professionals we are expected to hold ourselves accountable to the highest standards, and that starts with having tough conversations that really challenge the way we do business. That is very easy to say – but frankly much, much harder to do. Our natural biases will not want to confront the truth. As I explained in my article — "Thinking Slow on Cybersecurity" — the majority of time we will just do fast thinking — put together limited information to tell a coherent story we like. We fail to look for, or dismiss, any information that is contrary to our fast thinking story. I also explained Professor Daniel Kahneman's concept of Cognitive Ease, which is basically the more often we tell ourselves a story the greater confidence we have in that belief.

We have a battle going on. On one side is human nature. Our natural bias for fast thinking, to keep believing what we believe based on limited information, along with arrogance, fear and

embarrassment, which can keep us from facing facts and indications that say things aren't going as well as we would like.

On the other side of the equation is slow thinking, our ability to think critically about an issue, bringing all the best information to light, to overcome ignorance and arrogance. Along with slow thinking is our moral courage, which has to be strong enough to overcome fear and embarrassment.

Admiral Rickover established an amazing Navy Nuclear Power program that was based on technical excellence, rigor, and critical thinking. It's that ability to think critically, coupled with the ideas I have gleaned from a number of great books – most importantly — Thinking, Fast and Slow by Daniel Kahneman, that I hope to apply to the challenging problem of cybersecurity.

Last, I explained Kahneman's idea of What You See Is All There Is (WYSIATI) in my article. Too often people take limited information and quickly reach conclusions. If your readers aren't careful they may use WYSIATI to assume that I don't think the Navy and the DoD are making progress in cybersecurity. That simply isn't true. The Navy and DoD are making progress, and senior leadership is definitely engaged in this very important area. However, what I hope to do is share thoughts and examples that might help others to question their very thinking on these important issues and question whether we are getting the best value out of some of our efforts. To do that I will focus more on concerns and areas where I think we can make more progress, instead of sharing successes. With that let's get started.

Cyber value. We need to think hard about the idea of cyber value. What are we getting for the dollars we spend on cybersecurity? The DoD has been making progress for many years in bringing needed capabilities to our warfighters for a fair price. However, in an area where requirements are not clear, it would be easy to not only lose progress, but for things to get much worse.

Cybersecurity is one area where it can be a challenge to set clear requirements. How do you measure cybersecurity? How do you know if you have enough? How do you know how much you should spend on particular cybersecurity efforts? If you build in all sorts of security capabilities in a system that is low risk and low impact, you might pat yourself on the back when nothing happens, but what you will have missed is what you could have done with the time, money and resources that were wasted on a system that didn't need it. These are tough questions, and there is a high potential for waste.

Some people believe we should take a tactical approach to cybersecurity and fix every vulnerability to address cyber threats. You hear advocates of this approach say, "What affects one affects all!" and, "Our enemies only need to find one way in, we need to protect them all." For example, there is a big push today to eradicate older Windows operating systems and transition to the WIN 10 operating system. That makes sense for systems with high internet connectivity that will benefit from the new web browsing, cloud, and email security improvements that have been made.

However, there are thousands of disconnected, standalone systems throughout the Navy and DoD, that, while they use Windows operating systems, they don't use it for the applications mentioned, and thus would receive little benefit in security. We currently don't have the resources required to transition all the older Windows operating systems to WIN 10, so we need to think carefully about which systems will benefit most by making this transition. From a hacker standpoint — why would a hacker want to take the time to figure out how to gain access and attack a single standalone system, with limited means of propagating an attack, when there are systems available with many points of entry, and massive ability to propagate the attack?

We must continually remind ourselves that vulnerabilities by themselves are not risks. Risk assessments need to evaluate so much more, including: where vulnerabilities exist, the architecture and protections that surround them, the manner in which they are exploited, the capabilities adversaries need to exploit them, and also the benefit adversaries get by exploiting them. We recognize that this sort of assessment makes sense when we talk about our "defense-in-depth" strategy.

In order to get the best cyber value, the answer is rarely "one size fits all." It is more often characterized by "it depends." Yet when we talk about the "eradication of Windows XP" we can easily fail to recognize the real risks we are dealing with, which can be small, or the little value we are getting from some of these efforts.

In the Submarine Force, we train on the difference between a seawater leak and flooding. A seawater leak does not have the ability to sink our submarine, but flooding certainly does. Right now, we have many people worrying about the potential for every cyber seawater leak, instead of focusing our efforts on the cyber flooding scenarios that could result in the loss of all hands.

In contrast to the "fix every vulnerability" approach, there has been some great work done in the field of safety for many years concerning how to constrain system performance while ensuring system services. Professor Nancy Leveson, from the Massachusetts Institute of Technology, is one of the leaders in this field. In their paper, Systems Thinking for Safety and Security, Professor Leveson and Professor William Young explain,

> "The fundamental challenge facing security professionals is preventing losses, be they operational, financial or mission losses. As a result, one could argue that security professionals share this challenge with safety professionals. Despite their shared challenge, there is little evidence that recent advances that enable one community to better prevent losses have been shared with the other for possible implementation."

They go on to explain,

> "Cyber security has largely been framed as a tactics problem, focusing on how best to defend networks and other information assets against threats. While necessary, we believe this misses the greater objective of securing the systems' ability to produce the services and functions society depends on. Defending networks is not an end in itself; rather it is a means to protecting these higher-

> level services and missions against disruptions. "Reframing the problem into one of strategy may ultimately produce better outcomes. In practice, this reframing involves shifting the majority of security analysis away from guarding against attacks (tactics) and more towards the broader socio-technical vulnerabilities that allow disruptions to propagate throughout the system (strategy). Put another way, rather than focusing the majority of the security efforts on threats from adversary action, which are beyond the control of the security specialist, security efforts should be focused on the larger, more inclusive goal of controlling system vulnerabilities."

Let me provide an example of the difference between these two approaches that will help make clear where we are getting good cyber value and where we aren't.

Let's say you had an industrial control system like a shipboard diesel generator. One potential cyber concern is that if adversaries take control of the diesel generator control system they could overspeed the diesel generator, causing it to explode and damage the ship. The tactical approach would be to fix every vulnerability that has the possibility of causing that casualty.

Proponents of the tactical approach might take all of the following actions: Establish a Trusted Foundry to ensure computer hardware is free of malware; use secure software coding practices to reduce software defects; scan portable equipment used for maintenance for malware; control even unclassified information about the control system so that adversaries can't exploit it; disconnect the diesel generator control system from the ship's network so there is no network path to reach the control system; and have two watchstanders on duty at all times so that an insider threat can't implant malware.

I could go on and on trying to address potential vulnerabilities. The list of potential vulnerabilities is endless and the benefit is minimal, but the costs are not. There is not a lot of cyber value in this type of approach.

Now let's compare a more strategic approach on the lines of what Professor Leveson is advocating. In this instance, the risk we are concerned about is overspeeding the diesel generator. In the Navy, diesel generators have overspeed protection. This protection is often provided by a mechanical flywheel. When the speed of the diesel generator gets too high, the mechanical flywheel will trip, causing the diesel fuel supply to be cut off, thus slowing the diesel generator. With this design, there is no way to overspeed a Navy diesel generator via a cyber-attack.

All the vulnerabilities I mentioned above cannot be used to overspeed a diesel generator with this protection, because that system risk is appropriately constrained. This example demonstrates one simple approach to ensure that the system's behavior is appropriately constrained instead of trying the tactical approach of fixing every vulnerability.

The good news is we have been doing this sort of approach in the Submarine Force and also Naval Aviation for many years. Our naval nuclear reactors were designed so that they remain safe even in the face of personnel error, or mechanical and electrical faults. They remain safe because of the combination of socio (procedures and training) and technical (mechanical and

electrical design) constraints placed on the system, which have allowed naval nuclear reactors to steam safely over 70 million miles without a significant incident.

Fortunately, these same system constraints will also provide protection against cyber-attacks. This isn't bravado; it's good engineering. Similar design was put into the Naval Aviation's Air Worthiness Program to ensure planes don't crash just because of an electrical fault or a lightning strike. These types of approaches built resilient systems to ensure the most important services remained and the greatest concerns were addressed.

We need to move away from the tactical approach to cybersecurity and take a more strategic approach to constrain system performance while providing key services. That's how we will get better cyber value.

Q: In our last discussion, you likened the current cybersecurity response to "a lot of people pushing a lot of paper, very slowly." What technical, operational and process solutions would you recommend? Do you think a decentralized approach to cybersecurity would work better than the current model?

A: Let me start with the second question first. Our last Chief of Naval Operations, Admiral Jonathan Greenert, used to say we need to "Get Faster." I was at the Council of Foreign Relations many months ago when the Army Chief of Staff, General Milley spoke, and he said "Speed Matters." Our current CNO, Admiral John Richardson says, "We have got to change the way we do business in a fundamental way to keep up with this rate of introduction of technology and opportunity or we will just fundamentally … fall behind."

I think most people would agree that cybersecurity is an extremely fast problem. So if you want to go fast you need to do two things: No.1. You need to have more people doing. So yes, you must have a decentralized approach. No. 2. You have to be able to take risks. This is what we are doing a very poor job of in the world of cybersecurity today. When we should be doing more research, development, testing and evaluation — we are doing less because of the slow, centralized, certification and accreditation process we have in place.

We have plenty of networks that we would like to add new cybersecurity tools to, including continuous monitoring, intrusion detection and intrusion prevention systems tools, but we are so worried about doing something wrong that we spend months, and sometimes years, getting permission to connect the equipment and capabilities that we need to provide sound cybersecurity capabilities.

Franklin Covey has a great book and leadership course called The Speed of Trust. This isn't about blindly trusting people to do the right thing, but instead providing smart trust to the people who have the right training and qualifications and then giving them the authority to act. We need real capability now and we have people that are knowledgeable and prudent that could provide us those capabilities quickly. Unfortunately, we have paralyzed ourselves with paperwork that quite often results in no change to the system in question. Let me provide an example.

The DoD instruction for Public Key Infrastructure (PKI) and Public Key (PK) Enabling specifically states that standalone (disconnected) systems are not required to implement PKI. That makes sense because the purpose of PKI is to ensure the secure transfer of information across networks, so only the right people can get access to the information. If the system you are using is standalone it is not transferring information, and thus the risk of improper access to information is greatly reduced, so PKI is not necessary.

Unfortunately, even though the DoD instruction specifically states PKI is not required for standalone systems, the Navy policy is that a waiver request is still required for standalone systems not employing PKI. These waivers usually involve about three to four people to review, including a Flag Officer or a member of the Senior Executive Service, before it is sent to the Systems Command for review, who in turn must send it to the OPNAV staff for review.

Just recently at NAVSEA, we had 45 waiver requests for systems that are not implementing PKI, one third of those requests were for standalone systems. So here is an example where we could cut waste of resources by a third, waste that is impacting senior leadership and personnel across three organizations, by eliminating paperwork for something we aren't required to do and aren't going to do.

If we want to get faster, we need to take risks, and we need to recognize and eliminate processes that add little to no value. We need to ensure we have good guidance, and appropriately trained people, and then trust them to take the right action. So yes, I definitely believe we need to shift to a more decentralized approach.

As far as recommendations for changes to technical, operational and process solutions what I will offer are questions we should ask ourselves, along with some examples and recommendations where I think we can do better. I would categorize the questions into two groups. The first group of questions concerns understanding the return on investment (ROI) for our efforts, and the second group of questions deals with establishing conditions for success (CFS).

Concerning the ROI questions, for all our technical, operational and process initiatives, we shouldn't ask ourselves can we do these efforts, but should we be doing them? How do we know our initiatives are working? Can results be measured in a meaningful way? We need to evaluate not only what is the return on investment we are getting out of the effort, but what are the opportunity costs of the chosen approach?

The CFS questions are based on establishing conditions that are more likely to generate desired results. Are our initiatives helping us to learn? Are our efforts leveraging the talent and knowledge of our workforce, or are we just telling people what to do? Are we establishing clear priorities that focus on what is most important? Are our initiatives becoming more complex and more likely to fail, or are we setting them up with human frailties in mind, making them simpler, and more likely to succeed? Let me give some examples in regards to these questions.

For the ROI questions of, "How do we know our initiatives are working?" and, "Are we measuring the right things?" I will use the example of our C&A process. I would say what we

care about is not a piece of paper that says we have the Authority to Operate, but did anything meaningful change in terms of the security of the system as it went through the C&A process? Did the process cause any significant cybersecurity capabilities to be added which reduced the risk of the system? Were those changes made to the system in question, or did we just get a Plan of Action and Memorandum saying we promise we will make changes in the future?

If we measure actual cybersecurity changes made, and when those changes occurred, then we can get a sense of where value is being added in the process. If action is being performed by the system owners, then we could greatly reduce or eliminate the substantial number of reviews from outside organizations, reducing overall costs and improving cyber value. If system owners aren't taking appropriate actions to reduce cyber risks, then we can evaluate why that isn't happening.

In this manner, we would work to get those closest to the system, the system owners, to implement necessary cybersecurity improvements as quickly as possible. To be clear, I am not criticizing the people involved with certification and accreditation, I am questioning whether that process is serving us well. Are we getting good value from the C&A process, in the most efficient way, resulting in improved cybersecurity for our systems?

If results are not measureable we need to question if there is any value to the effort. One recent concern is the accessibility of unclassified but sensitive information. The idea here is that some information may not be classified but we still need to control it because if our adversaries get a hold of it then they will be able to use that information to exploit our systems.

It is possible to place additional constraints to control this unclassified information. These constraints will have a monetary cost and impact on our ability to move quickly. But how can you tell that these efforts are having any positive impact? There is no way to measure the supposed gains here. We have no idea how often adversaries are getting this information and putting it to use, so we can't even tell if measures put in place to reduce those events from happening are helping or not. Worse still, we already have in place requirements to classify information, so now we are just heaping bureaucracy upon bureaucracy with no clear return on investment.

Additionally, in situations like this, we often miss the opportunity costs that this sort of approach may have. For example, we talk quite often about how important small businesses are to the defense industry's ability to be innovative and efficient, and yet if additional requirements regarding controlling information are put in place, it is likely that will limit small business participation in the defense industry.

I don't have a great solution to offer here and I admit that it is very difficult to try to prove the value of a negative. Possibly the insurance industry can help us in this regard. That said, if there is no way to measure a positive outcome regarding a proposed initiative, I would say we should be cautious about implementing it.

The idea of understanding ROI must also be brought into play as we attempt to embrace innovation. I am certainly for trying new things and looking for better ways of doing business.

But before we mandate a new program across the Navy or DoD we should have some proof that it will yield positive results.

In his book, The Lean Startup — Eric Ries, advises that whenever you try something new, you need to prove as quickly as possible if your idea will actually work. Rapid prototyping is one example of getting some initial results that could predict your initiative will succeed. You don't want to make major investments to implement your plan and then find out there isn't a positive ROI.

The Navy is currently in the process of implementing the CYBERSAFE program. The idea behind CYBERSAFE is to provide additional security controls for a select number of critical systems. We should be clear that presently we do not know whether these controls will generate a positive return on investment. At this point, we don't have a good understanding what the additional actions will cost, nor do we have any evaluation of what improvements in security will be generated by the CYBERSAFE measures, compared with the cybersecurity actions already proscribed by programs, such as the Risk Management Framework. Will the gain in security be worth the investment in these additional cybersecurity measures?

With regards to establishing conditions for success, we should evaluate if we are putting in place conditions that make it more likely for us to succeed. For example, in our certification and accreditation process we have validators who are supposed to independently validate the certification and accreditation packages for systems they have been assigned. There is one big problem with this. The validators are paid for by the same programs they are preparing the certification and accreditation packages for. In doing this, we haven't put them in a good position to be independent.

What most programs really want is the Authority to Operate (ATO) so they can keep operating. As soon as the ATO is received, the validator no longer worries about that system and package, and moves on to the next package to be validated. The focus isn't on improving security for the systems. The validators are not rewarded for the amount of security problems they uncover and improvements in security they enable; rather, they are told "good job" when the program gets the piece of paper — the ATO. Instead of establishing conditions for success we have set them, and their program, up to fail.

One CFS would be does the initiative improve the opportunity for learning? I recently returned from a trip meeting with a number of ships' crews and commanders. One concern that was mentioned by one of the commanding officers was the limited visibility there is in terms of software patching. He informed me that in some cases software downloads were taking a long time to perform and yet he had no way of knowing the priority of the patches. In one instance, the ship terminated the download because it was limiting the ability to use a communications circuit. This is a perfect example where we aren't helping our people to learn and understand what we are asking them to do.

Not all software patches are created equal. If we treat them as such we may miss software patches that are extremely important, or we may limit a ship's operations for patching software that provides little to no value. Our cybersecurity efforts need to make crews more

knowledgeable instead of just telling them what to do. If we arm them with knowledge, they will figure out a good solution. If they are only doing what we tell them to do, then they will waste time and effort, and when it really matters they will be unprepared to take action.

We need to set clear priorities. It's clear the migration to Windows 10 is a priority for the DoD to improve the security of networks and information systems. The current direction is to transition "all computing systems currently running Windows operating systems… to the maximum extent practicable." As previously stated, we do not have the necessary resources to transition all DoD systems to WIN 10 in the timeline mandated by the DoD. When the necessary resources are not available, a better approach would be to devote resources to the highest risk, highest priority systems first. That is basically the concept behind the Risk Management Framework. Yet that is not the current approach.

The questions that are being asked are, "What is your plan to eradicate Windows XP and transition all systems to WIN 10?" So besides trying to figure out the plan for the high priority, high risk systems, there is also significant work being done on very low risk, low priority, disconnected systems.

Let me give an example of how we could change the current WIN10 implementation guidance to clearly set priorities. The guidance could be, "WIN 10 will be implemented to the maximum extent practicable for internet connected systems. Programs or commands intending to transition disconnected systems to WIN 10 must request permission to do so." By taking this approach we would find out the cost of transitioning the high priority, high risk systems to WIN 10.

We would also identify funds that might be shifted from lower priority systems to the high priority, high risk systems. Additionally, it would eliminate all the discussions between system owners, who do not want to transition disconnected systems because they believe there is little gain in security for the cost of implementation, with those above them who are more concerned with following DoD guidance.

In this manner we are clear about our priorities, focus our efforts on the highest risk systems, and eliminate distraction and confusion by removing all low priority, disconnected systems from the discussion. If it turns out we have resources to do all the highest priority systems, then lower priority systems could be addressed if deemed necessary.

Now let me refer back to the book Thinking, Fast and Slow and tell you how our natural biases will try to stop us from doing what I suggested above. The book explains the concept of "loss aversion" which is,

> "When directly compared or weighted against each other, losses loom larger than gains. This asymmetry between the power of positive and negative expectations or experiences has an evolutionary history. Organisms that treat threats as more urgent than opportunities have a better chance to survive and reproduce."

We hate losses, and are more inclined to worry about them than potential gains. When faced with the idea of focusing our efforts on just high priority systems, our DNA screams out — "Don't

forget those low priority systems, hackers could attack them too." What we don't do is recognize the gain we will make in progress on the high risk, high priority systems by significantly reducing effort on the low risk, low priority systems. Peter Drucker, the management guru, once said, "If you want something new, you have to stop doing something old." If we want to make the most progress on our high risk systems, we need to stop wasting time on our low risk systems.

Unfortunately, Fast Thinking is also working against us here. "The confirmatory bias of System 1 (Fast Thinking) favors uncritical acceptance of suggestions and exaggeration of the likelihood of extreme and improbable events." So when I say we shouldn't focus on low priority, low risk, disconnected systems, instead of agreeing that they are lower risk, someone might say, "What about Stuxnet, that virus was used to attack a disconnected system?" Yes that is true. It is possible to attack a "disconnected system." However, in the Stuxnet situation the reward for expending the resources to develop that attack was huge, the destruction of the Iranian nuclear centrifuges.

We have thousands of standalone systems that are used for various mundane operations, from allowing microscopic observations to be recorded on a laptop, to operating an industrial lathe. The majority of our standalone systems are not known to outside observers. Few, if any, of these standalone, disconnected systems, have the potential impact, or potential gain from an adversarial standpoint, of reaching the level of magnitude of the Iranian centrifuges. We must recognize that if the potential gain is low, and access is significantly limited, then the probability of attacking that type of disconnected system is also low.

We are making progress in our efforts regarding cybersecurity. We would do better if we thought more critically on the return on investment of our initiatives, and worked to establish conditions for success whenever possible.

Q: With the high velocity of ongoing cyber-threats, most organizations are mired in a continuous cycle of remediation and software patching. How can an organization get ahead of the curve?

A: First, I do think a continuous cycle of software patching makes sense for systems that have continuous connectivity. When patching makes sense, I am all for it. However, I think a much better approach that will actually get people ahead of the curve relates back to the ideas I mentioned earlier regarding Professor Leveson's work in the field of safety. How do we constrain system behavior to prevent the thing we are most worried about? Instead of having to fix every software flaw with a patch, what can be done so the bad thing doesn't happen?

Modern credit cards with chip technology are an example of this. Credit cards with chips create a one-time transaction code that can only be used once, unlike the magnetic strip credit cards that can easily be duplicated and reused. The chip constrained the risk of stealing the credit card information and using it over and over.

The other way we get ahead of the patching curve is to think about what we can do operationally to mitigate the risks we are concerned about. Here is one simple example of this type of approach. If you were concerned about a cyber-attack taking over a ship's rudder during an underway replenishment and causing a collision, you could analyze the attack vector and take simple steps to defeat it. You might secure the types of communications that adversaries might need to exploit to take over the steering system. You might shift steering control into a manual mode of operation that is not influenced by networks.

I am not advocating these changes, only giving them as food for thought for how process changes might help address cyber concerns, instead of implementing a technical solution. The great thing about these processes is that they could be used, or not, depending on the perceived level of risk.

Q: To address the ongoing cybersecurity problem in the federal government and Defense Department, some cybersecurity experts advise embracing a hacker culture with the active recruitment of hackers into the federal government and DoD. Others say hackers are social misfits, at best, and at worst, criminals. What are your thoughts on the kind of cybersecurity culture and workforce needed?

A: What hackers have always had, and what we need more of, is practical knowledge — knowledge not only of how our networks and information systems work, but knowledge about how people work. The "socio-technical" comment I brought up previously from Professor Leveson is exactly what I am talking about here.

Kevin Mitnick, one of the most famous hackers of all time, liked to call himself a "social engineer." It doesn't make sense for hackers to spend their time trying to break through extremely powerful encryption and other cyber technology if it is so much easier to take advantage of human flaws. Hackers aren't going to choose the Indiana Jones whip to fight with a tough encryption device, when they know they can use their spear phishing gun to make things easy. Hackers know what works when it comes to attacking our systems.

The Navy and DoD are doing more and more in the area of Red Teams, which allow DoD and Navy hackers to attack network and information systems to identify vulnerabilities. I really think this is one of our best successes and worst failures all at the same time. It's a great success because it just makes sense. If you want to do better against cyber-attacks, then you need to get into the field of battle and learn how that game is played. If we liken this to sports, why not play the game more often and start learning how to better defend our networks?

I really have great respect for the hacker community and I am happy to see that Navy and DoD leadership is pushing more of these activities. I attended a debrief from a Navy Red Team a number of weeks ago and the comment from the Red Team leader was, "It was all high fives and champagne for the Red Team." Meaning: once again the Red Team was able to easily attack the information system they were asked to attack.

Unfortunately, here is where we are completely broken when it comes to Red Teams and hacking. When Red Teams conduct exercises on a particular information system or network, their results, and the vulnerabilities they uncover, are only released to the system owner. Think about that for a moment. All of us that have information systems and networks are in this cyber-game whether we like it or not. And by the success of the Red Teams, we are showing we aren't very good at this game. So what do coaches do when their teams aren't doing well? They study the game. They watch the game. They read about the game. They share lessons learned with other coaches, and they play the game. But in the Navy and the DoD that is the exception instead of the rule, because lessons learned aren't shared unless the system owner allows that information to be released.

Millions of dollars are spent on Red Team activities, but we are failing to maximize the value we are getting out of those efforts. In fact, if it is a really great lesson learned that exposes a really big vulnerability, is the system owner incentivized to share that with others? Absolutely not. Which is why the rules are set up the way they currently are. I discussed this issue with Brad Horton, the director of the Threat System Management Office Red Team, one of the top Red Teams in all of DoD. Here is what he had to say:

> "As a Red Team, we often meet with system designers, Program Executive Officers, Program Managers, and others to share what we know. Without discussing the actual targets, though, we lose the impact and 'so what factor' necessary to affect needed changes. What we have is a culture problem. Commanders and Program Managers (target owners) are judged by their ability to survive cyber-attacks. Red Teams are threats to their promotions, cost, schedule, and the things that matter in the DoD. No commander ever received an award for letting a Red Team (a) attack them and find their weaknesses or (b) voice those results to the community. That, to me, is what has to change."

Many people believe that someday there will be a breach of a DoD information system with traumatic consequences. Unfortunately, when we do the post-incident assessment we will find that the vulnerabilities that are identified are things our Red Teams know today. We just never put in place a process to share this important information across the DoD.

This same problem with sharing lessons learned also exists with our cyber test ranges. The DoD Cyber Range in Quantico and the National Cyber Range in Orlando both have similar agreements with users that cyber test range results are only released to system owners. Again, I think these test ranges are wonderful tools. But failing to share what we are learning throughout the DoD is a tremendous waste of resources.

In most cases, we should be able to sanitize the results to get out important lessons learned, while still protecting the identity of the program. In cases where we can't protect the identity of the program, the good of the many should outweigh the good of the few, or the one. Trying to address our cyber challenges one program at a time is not a sound strategy. So as far as workforce and culture is concerned, I think we need to ensure we are more concerned about the good of the entire DoD over any one particular information system, network, or combat system.

Q: Phishing schemes are becoming increasingly sophisticated and tough to spot. Do you think personnel should be held accountable for cybersecurity breaches?

A: Somewhat. I really like the idea of actively spear phishing our own employees. In the past month and a half we have spear phished 3,000 NAVSEA employees. We are just starting this program and it is raising awareness for our employees to the spear phishing threat. We began this initiative, because many months ago I met with a major defense contractor who explained how they approached the spear phishing threat. Everyone in their company is spear phished at least once a year. The first time you fail a spear phishing attack you are notified. The second time you fail, you and your supervisor are called in to discuss why you repeatedly fail spear phishing attacks. The third failure is not an automatic termination, but it is grounds for dismissal.

This makes perfect sense. If someone is being careless do we want that type of person increasing the risk of a successful cyber-attack for our networks? If employees fall prey to attacks that the majority of employees were able to recognize, I think there needs to be some action taken, not necessarily punitive, unless they repeatedly show they are careless. However, in some circumstances I think sophisticated adversaries will be able to trick even well trained and alert employees. In those cases we should recognize that the attack was sophisticated beyond the employee's ability to detect it.

Q: Is there anything else you would like to discuss?

A: Absolutely. I'd like to apologize for the length of this online interview. Not in pages, but in the amount of time it took me to provide my answers. It's been over two months since you first gave me the questions and I have tried to set a good example of slow thinking by challenging the ideas that I've shared with you.

I have shared these ideas with representatives from the OPNAV staff, SPAWAR, and FLTCYBERCOM. I have requested feedback from experts in the field of safety, Red teams, cyber ranges and Certification and Accreditation. Having this sort of open dialogue is what is necessary to make progress on challenging issues.

The majority of the feedback I received confirmed my concerns. I was even fortunate enough to get feedback from the Nobel Prize winning author and professor, Daniel Kahneman, who I am happy to report felt that "Thinking Slow on Cybersecurity" did an "excellent" job of applying the concepts from his book, "Thinking, Fast and Slow."

However, in some cases, based on the feedback, I discovered I was stuck in my own bit of fast thinking. In one area in particular, Certification and Accreditation, I realized I was unfairly dismissing some of the ways that process brings value to the DoD. I updated what was written in this interview to reflect what I now believe is a more accurate assessment of our current state of affairs in regards to cybersecurity.

A few people I asked for feedback were concerned with the idea of raising issues that could possibly reflect poorly on Navy and DoD initiatives. You probably won't be surprised that I don't see it that way.

When we speak of our values in the Navy and DoD, what really matters is not what we say or the appearance of how we are doing. What matters is if our actions reflect our values. Being able to think slowly, and critically, and engage in tough discussions such as this, holding ourselves accountable for the safety and security of our systems and our people, and having the integrity and moral courage to openly address some of our greatest cybersecurity challenges, are all sure signs that we are committed to the values of Honor, Courage, and Commitment we proclaim. I think that reflects very positively on the Navy and DoD, and it is exactly what we will need to effectively address the cybersecurity challenges that are ahead of us.

Thanks again for the opportunity to discuss these important issues with you.

Captain Zimmerman commanded USS JEFFERSON CITY (SSN 759) from 2006-2008. From 2012-2015 he was the Major Program Manager for the Submarine Combat and Weapon Control Systems Program (PMS 425). The PMS 425 program received the 2014 AFEI 2014 Government Award for Innovation, Cost Savings, and Operational Impact, and the 2015 NAVSEA Commander's Team Innovation Award for the SSBN Tactical Control System Upgrade. Capt. Zimmerman was recognized by the Secretary of the Navy for the 2015 Innovative Leadership Award (Honorable Mention).

** = The opinions expressed here are solely those of the author, and do not necessarily reflect those of the Department of the Navy, Department of Defense or the United States government.